

Patent Retrieval: A Literature Review

Walid Shalaby, Wlodek Zadrozny
Department of Computer Science
UNC Charlotte
{wshalaby, wzadrozny}@uncc.edu

ABSTRACT

With the ever increasing number of filed patent applications every year, the need for effective and efficient systems for managing such tremendous amounts of data becomes inevitably important. Patent Retrieval (PR) is considered is the pillar of almost all patent analysis tasks. PR is a subfield of Information Retrieval (IR) which is concerned with developing techniques and methods that effectively and efficiently retrieve relevant patent documents in response to a given search request. In this paper we present a comprehensive review on PR methods and approaches. It is clear that, recent successes and maturity in IR applications such as Web search can not be transferred directly to PR without deliberate domain adaptation and customization. Furthermore, state-of-the-art performance in automatic PR is still around average. These observations motivates the need for interactive search tools which provide cognitive assistance to patent professionals with minimal effort. These tools must also be developed in hand with patent professionals considering their practices and expectations. We additionally touch on related tasks to PR such as patent valuation, litigation, licensing, and highlight potential opportunities and open directions for computational scientists in these domains.

1. INTRODUCTION

Patents represent proxies for economic, technological, and even social activities. The Intellectual Property (IP) system motivates the disclosure of novel technologies and ideas by granting inventors exclusive monopoly rights on the economic value of their inventions. Patents, therefore, have a major impact on enterprises market value [1]. With the continuous rise in the number of filed patent applications every year, the need for effective and efficient systems for managing such tremendous amounts of data becomes inevitably important.

Typical patent analysis tasks include: 1) technology exploration in order to capture new and trendy technologies in a specific domain, and subsequently using them to create new innovative services, 2) technology landscape analysis in order to assess the density of patent filings of specific technology, and subsequently direct R&D activities accordingly, 3) competitive analysis and benchmarking

in order to identify strengths and differences of corporate's own patent portfolio compared to other key players working on related technologies, 4) patent ranking and scoring in order to quantify the strength of the claims of an existing or a new patent, and 5) prior art search in order to retrieve patent documents and other scientific publications relevant to a new patent application. All those patent-related activities require tremendous level of domain expertise which, even if available, must be integrated with highly sophisticated and intelligent analytics that provide cognitive and interactive assistance to the users.

Patent Retrieval (PR) is the pillar of almost all patent analysis tasks. PR is a subfield of Information Retrieval (IR) which is concerned with developing techniques and methods that effectively and efficiently retrieve relevant patent documents in response to a given search request. Although the field of IR has received huge advances from decades of research and development, research in PR is relatively newer and more challenging. On the one hand, patents are multi-page, multi-modal, multi-language, semi-structured, and metadata rich documents. On the other hand, patent queries can be a complete multi-page patent application. These unique features make traditional IR methods used for Web or ad hoc search inappropriate or at least of limited applicability in PR.

Moreover, patent data is multi-modal and heterogeneous. As indicated by Lupu et al. [2], analyzing such data is a challenging task for many reasons; patent documents are lengthy with highly complex and domain specific terminology. To establish their work novelty, inventors tend to use jargon and complex vocabulary to refer to the same concepts. They also use vague and abstract terms in order to broaden the scope of their patent protection making the problem of patent analysis linguistically challenging.

PR starts with a search request (query) which often represents a patent application under novelty examination. Therefore, several methods for query reformulation (QRE) have been proposed in order to select, remove, or expand terms in the original query for improved retrieval. QRE methods are either keyword-based, semantic-based, or interactive. Keyword-based methods work by searching for exact matches between search query terms and the target corpus, and thus fail to retrieve relevant documents which use different vocabulary but have similar meaning to original query. In order to alleviate the vocabulary mismatch problem, semantic-based methods try to search by meaning through expanding queries and/or target corpus with similar or related terms and thus bridging the vocabulary gap. Because neither methods proved acceptable performance, few interactive methods were proposed to allow users to interactively control QRE with reasonable effort.

Table 1: Patent kind codes of major patent offices

Type	USPTO (US)	EPO (EP)	WIPO (WO)
A1	application	application w/ search report	
A2	republished application	application w/o search report	
A3	-	search report	
A4	-	supplementary search report	publication of amended claims
A9	modified application		
B1	granted patent w/o A1	granted patent (publication)	-
B2	granted patent w/ A1	amended B1	-

This review aims to provide researchers with an illustrative and critical overview of recent trends, challenges, and opportunities in PR. The rest of this paper is organized as follows. Section 2 presents some preliminaries and background about patents data. Section 3 provides an overview of evaluation tracks and data collections for PR benchmarking. An illustration of PR tasks is presented in Section 4. Section 5 presents a comprehensive review on PR methods and approaches. Section 6 lightly touches on related tasks such as patent quality assessment, litigation, and licensing. Finally, concluding remarks are presented in Section 7.

2. PRELIMINARIES

2.1 Patent Documents and Kind Codes

Patent documents are mostly textual. They are highly structured with typical elements (sections) including *title*, *abstract*, *description* (*a.k.a background of the invention*), and *claims*. The *description* section articulate in details the technical specification of the invention and its possible embodiments. The *claims* section is the most significant one as it describes the scope of protection sought by the inventor and hence encodes the real value of the patent. Patent documents are lengthy with highly complex and domain specific terminology. They also contain multiple data types (e.g., text, images, flow charts, formulae...etc) with a rich set of meta-data and bibliographic information (e.g., *classification codes*, *citations*, *inventors*, *assignee*, *filing/publication dates*, *addresses*, *examiners*...etc).

Typically, each patent has a set of pertaining documents which published throughout its life-cycle. All documents are identified by an alphanumeric name with a common naming convention. Names start with two letters identifying the issuing patent office (e.g., US and EP), then the patent number as sequence of digits, and finally a suffix indicating the document’s kind code. The kind code identifies the stage in the patent life-cycle at which the document is published. Table 1 shows a brief description of kind codes used at major patent offices including the US Patent and Trademark Office¹ (USPTO), the European Patent Office² (EPO), and the World Intellectual Property Organization³ (WIPO).

2.2 Patent Classification

Patent offices organize patents by assigning classification codes to each of them based on the technical features of the invention. The patent classification system is a hierarchical one. Common classification systems include the International Patent Classification (IPC), the US Patent Classification (USPC), and the Cooperative Patent Classification (CPC).

¹<http://www.uspto.gov/>

²<http://www.epo.org/>

³<http://www.wipo.int/portal/en/index.html>

2.3 Patent Families

A patent family is patent documents that refer to the same invention and are published by different patent offices around the world [3], usually in different languages depending on the issuing patent office. Patent families could be exploited to expand the prior art list of topic patents as they disclose the same invention.

3. DATA & EVALUATION TRACKS

This section presents an overview of evaluation tracks organized for patent data analysis along with available data collections with focus on tasks pertaining to PR.

3.1 CLEF-IP Collections

The Conference and Labs of the Evaluation Forum⁴ (CLEF) is an European series of workshops which started in 2001 to foster research in Cross Language Information Retrieval (CLIR). The Intellectual Property (IP) track (CLEF-IP) which ran between (2009-2013) was organized to: 1) foster research in patent data analysis, and 2) provide large and clean test collections of multi-language patent documents, specifically in the three main European languages (English, French, and German). Research labs have the opportunity to test their methods on multiple shared tasks such as PR, patent classification, image-based PR, image classification, flowchart recognition, and structure recognition [1, 3–6].

The CLEF-IP data collection are patent documents extracted from the EPO data. It is provided through the Information Research Facility⁵ (IRF) and hosted by Marec⁶. Patent documents are provided in XML format and have common Document Type Definition (DTD) schema. The collection was constructed according to the proposed methodology by Graf and Azzopardi [7] and is divided into two pools:

1. The corpus pool: documents selected from this pool are provided for participating labs as training or lookup instances depending on the task.
2. The topics pool: documents selected from this pool are called topics and they represent testing or evaluation instances depending on the task. For example, in prior art search, the topic might be a patent application document for which it is required to retrieve prior art.

The XML documents consist of the main textual sections such as *bibliographic data*, *abstract*, *description*, and *claims*. Each section

⁴<http://www.clef-initiative.eu/>

⁵<http://www.ir-facility.org/>

⁶<http://www.ir-facility.org/prototypes/marec>

is written in one or more languages (English, French, and/or German) and is denoted by a language code. At least the claims of granted patents (B1 documents) are written in the three languages because it is EPO requirement once a patent application is granted.

CLEF-IP 2009 Collection: this dataset was designed for the prior art search task [4]. The corpus pool contains documents published between (1985-2000) (~2m documents pertaining to ~1m unique patents). The topics pool contains documents published between (2001-2006) (~0.7m documents pertaining to ~0.5m individual patents). Topics are sets of documents from the topics pool with sizes ranging from 500 to 10,000 topics. Topics were assembled from granted patent documents including *abstract*, *description*, and *claims* sections. Citation information from the *bibliographic data* section was excluded.

A major pitfall in this dataset is its topics, which were chosen from granted patent documents (B1 documents). Initially, the creators of the dataset were motivated by having topics from granted patent documents which have claims in three languages. This was thought to provide a kind of parallel corpus suitable for CLIR. The problem of using such documents is simple, it contradicts the practice of IP search professionals who start with the patent application document not the granted one.

CLEF-IP 2010 Collection: this dataset was created for the prior art search and patent classification tasks [5]. The corpus pool of this dataset contains documents with publication date before 2002 (~2.6m documents pertaining to ~1.9m unique patents). The topics pool contains documents published between (2002-2009) (~0.8m documents pertaining to ~0.6m unique patents). Topics for the prior art task are two sets of documents from the topics pool; a small set of 500 topics and a larger set of 2000 topics. Unlike the CLEF-IP 2009 dataset, topics are assembled from patent application documents rather than granted patent documents.

CLEF-IP 2011 Collection: This dataset was created as a test collection for four tasks: prior art search, patent classification, image-based prior art search, and image classification [1]. The topics and corpus pools were the same as in CLEF-IP 2010 dataset. For the prior art task, 3973 topics were provided as a separate archive of patent application documents.

CLEF-IP 2012 Collection: this dataset was created as a test collection for three tasks: passage retrieval starting from claims, chemical structure recognition, and flowchart recognition [6]. The topics and corpus pools were the same as in CLEF-IP 2010 dataset. The passage retrieval task is designed differently from previous CLEF-IP prior art search collections. The purpose for this task is to retrieve both documents and passages relevant to a set of claims. Topics for the passage retrieval task were extracted from patent applications published after 2001. Relevance judgments were the highly relevant citations only (i.e., marked X or Y) in the examiners' search reports (A4 documents) of chosen topic patents.

CLEF-IP 2013 Collection: this dataset was created as a test collection for two tasks: 1) passage retrieval from claims, and 2) structure recognition from patent images [3]. The topics and corpus pools were the same as in CLEF-IP 2010 dataset. Similar to CLEF-IP 2012, the CLM task is designed to retrieve both documents and passages relevant to a set of claims. Topics for the passage retrieval task were extracted from patent applications published after 2002. Overall, the topics set contained 148 topics extracted from

69 patent applications.

3.2 NTCIR Collections

The Japanese National Institute of Informatics Testbeds and Community for Information access Research project⁷ (NTCIR) started in 1997 to support research in IR and other areas, focusing on CLIR. NTCIR has been organizing a series of workshops providing test collections to researchers for evaluating their methodologies on multiple CLIR tasks [8]. Between NTCIR-3 and NTCIR-11 (2002-2013), there has been dedicated tasks for patent data analysis including patent retrieval [9], classification, mining, and translation.

NTCIR-3: the PR task in NTCIR-3 targeted the "technology survey" problem. The dataset for this task includes: 1) full text of Japanese patent applications between (1998-1999), 2) *abstract* of Japanese patent applications between (1995-1999) along with their respective English translations, and 3) 30 search topics where each topic includes a related newspaper article. The task is to retrieve patents relevant to news articles. Both cross-genre experiments in which patents were retrieved by a newspaper clip as well as ordinary ad hoc retrieval of patents by topics were conducted [9].

NTCIR-4: two PR tasks were organized in NTCIR-4 [10]: 1) patent map generation, and 2) invalidity search. The dataset for the PR tasks includes: 1) unexamined Japanese patent applications published between (1993-1997) along with English translations of the *abstract*, and 2) 34 search topics where each topic is a claim of a rejected patent application which was invalidated because of existing prior art. Relevance judgments were individual patents that can invalidate a topic claim by its own or in conjunction with other patents. Relevant passages to the invalidated claim were also annotated and added to the relevance judgments.

NTCIR-5: two PR tasks were organized in NTCIR-5 [11]: 1) document retrieval (invalidity search), and 2) patent passage retrieval. The dataset for the invalidity search task includes: 1) unexamined Japanese patent applications published between (1993-2002) along with English translations of the *abstract*, and 2) 1200 search topics where each topic is a claim of an invalidated patent application. Relevance judgments were generated in a manner similar to the one used in NTCIR-4 invalidity search task.

NTCIR-6: two PR tasks were organized in NTCIR-6 [12]: 1) Japanese retrieval (invalidity search), and 2) English retrieval. The dataset for the Japanese retrieval task is the same one used in NTCIR-5 but more topics were used (1685 topics). The English retrieval task was focusing on finding all the citations cited by the applicant and the examiner. The dataset for this task includes: 1) granted patents from the USPTO between (1993-2000), and 2) 3221 search topics where each topic is a granted patent published between (2000-2001).

3.3 TREC-CHEM Collections

The TREC-CHEM track was organized to motivate large scale research on chemical datasets, especially chemical patent IR [13].

TREC-CHEM 2009: this collection was created as a test collection for two tasks [13]: 1) technology survey, and 2) prior art search. 18 topics were provided for the technology survey task where relevance judgments were obtained from experts and chemistry graduate students. For the prior art search, 1,000 patents were

⁷<http://research.nii.ac.jp/ntcir>

Table 2: Scenarios of prior art search

Task	Who	When	Purpose	Output
related work search	inventor/prosecutor	pre-grant	all related work	applicant’s disclosure
patentability search	prosecutor/examiner	pre-grant/examination	novelty breaking work	grant/modify/reject
infringement search	owner/investor	post-grant	relevant claims/infringing products	sue/license/clearance
freedom to operate search	investor	post-grant	relevant claims/related work	clearance
invalidity search	competitor/defendant	post-grant	novelty breaking work	re-examine/IPR/PGR
technology survey	technology analyst	pre/post-grant	all published patents	survey report

provided as test topics where relevance judgments were collected from the citations of topic patents as well as their family members. The search corpus contains ~1.2m chemical patents filed until 2007 at EPO, USPTO, and WIPO. It also contains 59K scientific articles.

TREC-CHEM 2010: this collection was created for the same two tasks as in TREC-CHEM 2009 [14]. 30 topics were provided for the technology survey task. The search corpus contains ~1.3m chemical patents and 177K scientific articles. Relevance judgments were created the same way as in TREC-CHEM 2009.

TREC-CHEM 2011: this collection was created for the same two tasks as in previous TREC-CHEM tracks besides a new chemical image recognition task. The technology survey task topics were biomedical and pharmaceutical patents [15].

3.4 Other Sources

Other IP data sources are detailed in [16]. These include full patent texts as well as bibliographic information from major patent offices such as the USPTO, EPO, and WIPO. Bibliographic information for patents published from 1976 to 2006 is provided through the National Bureau of Economic Research⁸ (NBER) and subsequently cleaned and extended to include patents until 2013⁹. Patent prosecution histories are available through the Patent Application Information Retrieval¹⁰ (PAIR). Patent assignments, filings, classifications, and petition decisions are also provided through the USPTO bulk downloads hosted by Google¹¹.

4. PATENT RETRIEVAL TASKS

The goal of PR is to retrieve relevant patent documents to a given search request (query). This request can take different forms such as a sequence of keywords, a memo, or a complete text document (e.g. a patent application). The purpose of this task is manifold, for example:

- Retrieve related patents to a given patent application in order to gather related work, or invalidate one or more of its claims.
- Explore patent filing activity under specific technology.
- Explore the competitive landscape of a given company by looking at other companies filing patents similar to the given company patents.

Because of these multiple objectives, various PR tasks were proposed to fulfill each objective, and multiple datasets were provided depending on the given task.

⁸<https://sites.google.com/site/patentdataport/>

⁹<http://rosencrantz.berkeley.edu/batchsql/>

¹⁰<http://portal.uspto.gov/pair/PublicPair>

¹¹<https://www.google.com/googlebooks/uspto-patents.html>

Prior-art search is the main theme of the CLEF-IP and NTCIR tracks. The importance of this task stems from the requirement by all patent offices that filed patents must constitute novel, non-obvious, and non-abstract ideas. Therefore, an important activity through the patent life-cycle is to thoroughly ensure that no earlier published patent or material describing the prescribed ideas exist. The task can be defined as follows:

Problem: Given a patent application X, retrieve all related documents to X.

Prior art search is a total recall task, therefore it demonstrates several challenges. Search coverage is one of the main challenges, because it is required to cover all previously published material (patent or non-patent literature) in all forms (electronic or printed) which is infeasible. Another major challenge is the need to search through materials written in different languages. Last but not least, traditional IR methods perform poorly when confronted with the patent prior art search task. Mainly because the patent language is full of jargon and user defined terminology. Inventors intentionally tend to use different vocabulary to express same or similar ideas in order to establish the novelty of their work.

Prior art search is performed at different stages of the patent life-cycle, by different stakeholders, for various purposes, and for limited period of time. Understanding the real-life practices of patent professionals is critical to better satisfy their information need [17]. In other words, the search scenario depends on when it is done, by who, and for what reason(s). Table 2 shows these various scenarios which are detailed below.

Related Work Search: during the pre-grant stage, inventors and prosecutors run related work search to retrieve all relevant work to the invention. Moreover, some patent offices request from inventors an applicant’s disclosure document specifying all related publications when filing a new application.

Patentability Search: during the examination stage, patent examiners perform patentability search in order to ensure that the proposed ideas are novel, non-obvious, and non-abstract. The output of this task would be a search report with all retrieved relevant publications. In this report, each entry will have a special code indicating whether it is just a related publication, or novelty breaking one. Examiners would also specify which passages or figures in retrieved publications constitute relevancy. Depending on the search findings, the patent office might grant, reject, or ask the applicant to modify the patent application. Patentability search is also performed by patent prosecutors as a sanity check. Although this task should be of equal interest to prosecutors who file the patent application as it is to examiners, prosecutors often do not dig deep searching for relevant publications, and delegate finding relevant prior work to examiners in order to save costs.




Figure 1: Taxonomy of Patent Retrieval (PR) methods

Infringement Search: this task, also called product clearance search, aims to ensure whether an existing or a proposed product is infringing any published patent claim(s). Patent owners require that type search to find out if a third party has a product with features which are cited by one or more claims of their patents. If so, they might either sue or negotiate a license with that infringing party. The scope of the search here would include all related publications to commercial products (e.g., product descriptions, vouchers...etc).

Investors and R&D managers, on the other hand, require that type of search to ensure newly proposed product(s) are not infringing a published patent claim(s) and investment in such products would be lucrative. The scope of search in this case would be limited to patent and copyrighted literature only. Deep understanding and correct interpretation of patent claims is imperative for building the correct correspondence between product features and claims in order to establish or dismiss infringement.

Freedom to Operate Search: this PR task extends beyond infringement search. Here, investors and R&D managers not only need to make sure that proposed products do not infringe an existing patent or copyrighted material, but also to ensure they have the freedom to file patents on these products without worrying about previously prior art that might invalidate such inventions. Another objective of freedom to operate search is to make better investment decisions and R&D plans according to existing prior art.

Invalidity Search: as patents guarantee monopoly rights to their owners on the economic value of granted inventions, companies and other parties usually monitor granted patents of their competitors or pertaining to their technology landscape to ensure competitive superiority. Therefore, invalidity search is performed to find published material that was missed by the patent office during patentability search. Invalidity search is also considered as the first line of defense when a party is confronted with patent infringement lawsuit. Again published material might include patent or non-patent literature such as books, news articles, academic periodicals...etc. After finding such validity breaking material, a third party might file a post-grant (opposition) procedure depending on the patent office policies. For example, the USPTO provides procedures such as re-examination, Inter Partes Review¹² (IPR), and

Post Grant Review¹³ (PGR) in front of the Patent Trial and Appeal Board¹⁴ (PTAB).

Technology Survey: another PR task where, in a typical scenario, business managers would request search professionals to prepare a survey of patent documents given a memorandum they prepared from some source (e.g., news article) [9]. This basic scenario is limited only to patent literature and it is assumed that patent documents are just a collection of technical papers.

5. PATENT RETRIEVAL METHODS

In this section, we present a comprehensive review of PR methods and approaches. We start by presenting available test collections and evaluation metrics. Then, we provide a taxonomy of these approaches highlighting their characteristics and limitations.

As shown in Figure 1, PR methods can be categorized depending on which piece(s) of data from both the search queries and the search corpus are used for retrieving relevant documents. Keyword-based methods utilize only terms from search queries and look for exact matches in the target corpus. Pseudo Relevance Feedback methods utilize terms from the top ranked results of running the initial query to improve the set of relevant retrieved results. Semantic-based methods try to overcome the vocabulary mismatch problem between the search terms and related patents vocabulary by matching them based on their meanings. Metadata-based methods exploit the language independent non-textual metadata and bibliographic information in order to improve patent retrievability. Finally, interactive methods aim to better organize and present search results to the users. Moreover, through interaction, users are engaged in an iterative process of searching, reviewing, and refining hoping to retrieve as many relevant results as possible.

5.1 Test Collections & Evaluation Measures

As we highlighted in section 3, several datasets were created to support evaluating different PR techniques. In almost all of these datasets, relevant documents to search queries were collected from the citations of topic patent documents (e.g., CLEF-IP 2009/2010/2011 collections). Because these citations represent related prior work, they are appropriate only for the related work search task.

¹²<http://www.uspto.gov/patents-application-process/appealing-patent-decisions/trials/inter-partes-review>

¹³<http://www.uspto.gov/patents-application-process/appealing-patent-decisions/trials/post-grant-review>

¹⁴<https://ptabtrials.uspto.gov>

Table 3: Keyword-based Methods

Method	Description	Dataset	MAP	P	R	PRES
Verberne & E. D’hondt [18]	<ul style="list-style-type: none"> remove stopwords, and punctuation use claims as BOW 	CLEF-IP 2009	0.05	0.01	0.22	-
Magdy et al. [19]	<ul style="list-style-type: none"> remove stopwords, and frequent terms use different sections with manual weights perform IPC filtering use bigrams with $tf > 1$ 	CLEF-IP 2009	0.12	-	0.63	-
Mahdabi et al. [20]	<ul style="list-style-type: none"> use query language models on different sections use queries of 100 terms perform IPC filtering 	CLEF-IP 2010	0.12	-	0.60	0.49
Wang & Lin [21]	<ul style="list-style-type: none"> use linguistic-based concepts concept weighting using weighted tf-idf and mutual information 	CLEF-IP 2010	0.10	-	0.48	0.40
Konishi [22]	<ul style="list-style-type: none"> patterns to identify <i>claim</i> components terms patterns for explanation terms from <i>description</i> rank boosting based on IPC 	NTCIR-5	0.20	-	-	-

In other datasets such as CLEF-IP 2012/2013 collections, relevant documents were collected from novelty breaking citations found in examiners’ search reports, therefore, these datasets are appropriate for the patentability and invalidity search tasks. Though invalidity search requires non-patent literature as well.

Standard IR as well as PR-specific evaluation measures are generally used to evaluate PR systems including:

1. Precision (P) and Recall (R) at top- K ranks (e.g., $K=\{1, 5, 10, 50, 100, 1000\}$).
2. Mean Average Precision (MAP) [23] which generally favors early retrieval of relevant documents with less focus on recall.
3. Normalized Discounted Cumulative Gain (nDCG) [24] which favors not only early retrieval of relevant documents but also the respective ranking quality of these documents.
4. Patent Retrieval Evaluation Score (PRES) [25] which was proposed specifically for recall-oriented tasks such as PR. PRES focuses on the overall system recall as well as user’s review effort which can be estimated from the rankings at which relevant documents are retrieved.

5.2 Query Reformulation (QRE)

The most widely used techniques for PR are the Query Reformulation (QRE) techniques. These methods aim at transforming the input query Q into \bar{Q} by means of reduction or expansion of Q terms in order to improve the retrievability of relevant documents. QRE can be performed through:

- **Query Reduction (QR):** where a representative subset of terms are selected from Q and used as \bar{Q} terms. Position-based methods are the most commonly used in this category where terms from specific parts or sections of the patent document are used, or given higher matching weight than others. Another example of QR is the IPC-based methods which utilize terms from IPC definitions as a lexicon or stop-words list for Q .

- **Query Expansion (QE):** where representative terms other than the ones in Q are extracted and merged with Q to form \bar{Q} . Pseudo Relevance Feedback (PRF) methods are the most prominent in this category where terms from top ranked results of running Q are used to expand Q terms assuming these top results are relevant [26]. Other semantic-based QE methods work by expanding Q with terms of similar meanings such as synonyms or hyponyms.

- **Hybrid (QE & QR):** where irrelevant terms are removed from Q and more relevant terms are appended to Q to form \bar{Q} . Most techniques used for QE are appropriate for QR as well, where only terms appearing in the expansion list are kept and all others are pruned.

5.2.1 Keyword-based Methods

This set of techniques retrieves relevant documents by looking for exact matches between search query term(s) and the target data. keyword search operates under the closed vocabulary assumption where vocabulary is derived solely from terms that appear in the target search data. Table 3 shows some keyword-based methods along with their performance results on benchmark datasets. Keyword-based techniques differ in: 1) which elements of the target data are indexed, 2) which query terms are selected/removed, 3) the relative weights of such terms, and 4) the match scoring function.

Query Reduction (QR): the rationale behind QR approaches is intuitive as patents are very long documents with several sections. Querying with the whole document would be impractical and inefficient. Some QR methods are position-based; they select relevant terms based on their position in the patent document [18–20,31,32]. For example, Verberne and D’hondt [18] used only terms from the *claims* section on the CLEF-IP 2009 collection. However, the results were moderate in terms in MAP compared to other runs on the same collection.

Magdy et al. [19] experimented using text from different sections of the topic patent on the CLEF-IP 2009 collection. The authors used various combinations of sections including: 1) short sections such as *title*, *abstract*, first line of the *description*, first sentence of the *claims*, and 2) lengthy sections such as the *description* and

Table 4: Pseudo Relevance Feedback Methods

Method	Description	Dataset	MAP	PRES
Bouadjene et al. [27]	<ul style="list-style-type: none"> • use different methods of QE and QR from the PRF set • use Ricchio, MMR, LM 	<ul style="list-style-type: none"> • CLEF-IP 2010 • CLEF-IP 2011 	<ul style="list-style-type: none"> • 0.13 • 0.10 	<ul style="list-style-type: none"> • 0.55 • 0.45
Magdy et al. [19]	<ul style="list-style-type: none"> • naive PRF • remove stop-words • use most frequent terms 	CLEF-IP 2009	0.05	-
Mahdabi & Crestani [28]	<ul style="list-style-type: none"> • build regression model using relevance score, RF similarities...etc • use the model to estimate the effectiveness of RF • use top 100 RF and maximize AP 	CLEF-IP 2010	0.16	0.56
Ganguly et al. [29]	<ul style="list-style-type: none"> • perform query segmentation • retain segments highly to be generated using RF LM 	CLEF-IP 2010	0.14	0.47
Golestan Far et al. [30]	<ul style="list-style-type: none"> • manually annotate one relevant RF result • add terms in the annotated result to the query 	CLEF-IP 2010	0.29 ¹⁵	-
Golestan Far et al. [30]	<ul style="list-style-type: none"> • assume relevant RF results are known • add terms more frequent in relevant than irrelevant RF to query 	CLEF-IP 2010	0.48 ¹⁶	-

the *claims*. The authors assigned different weights to each section manually. Their best scores were achieved using a combination of all short sections and post-filtering retrieved documents keeping only those that share the same IPC classification code with the topic patent. The main challenge with such approach is how to assign the respective weight of each section automatically. Moreover, IPC filtering wouldn't be possible when only partial patent application is available for prior art search.

Mahdabi et al. [20] proposed a position-based QR method which selects relevant query terms by building two query language models using various sections of the topic patent: 1) a variant of the weighted log-likelihood model [33], and 2) a model based on the parsimonious language model [34]. Their experiments showed that queries constructed from terms in the *description* section using weighted log-likelihood give better results than other sections which agrees with the results in [27, 35, 36]. The main advantage of this approach is that, respective weights of query terms are derived automatically from the query model. However, some challenges still exist regarding tuning the model parameters such as the smoothing parameter which was set heuristically.

Query Expansion (QE): pattern-based QRE was proposed in many studies [22, 37, 38]. Wang and Lin [38] proposed patterns in the form of syntactic rules in order to extract query terms as weighted concepts. In [22], Konishi proposed a pattern-based QE method for the patent invalidity search task on the NTCIR-5 collection. Rather than using raw terms from topic patent's *claims* which are often abstract, the author, using pattern matching, identifies other specific terms in the *description* and use them as expansion terms. First, components of the invention are extracted from the topic *claim* using handcrafted patterns. Secondly, explanation sentences describing components of the invention are extracted from the *description* using handcrafted patterns. Thirdly, terms from first and second steps are used as the new query. The results showed that this QE approach works better than using terms extracted from the *claims* section only. The main drawback of this method is its dependency on manually coded patterns to identify potential terms. Meanwhile, it demonstrates the potential of using entities and their relations as retrieval features motivating the need for deeper and more generic

linguistic analysis of patent texts.

5.2.2 Pseudo Relevance Feedback (PRF)

These methods are one of the prominent techniques used for QRE. PRF starts with an initial run of the given query Q . Then, terms from top ranked results are used to select, remove, and/or expand terms in Q , assuming that these top results are relevant. PRF is thus advantageous as it works automatically without human intervention but might be computationally inefficient especially with long queries. Table 4 shows some PRF methods along with their performance results on benchmark datasets.

Despite their effectiveness and popularity, several challenges arise when it comes to PRF-based QRE [27] such as: 1) which part(s) of the patent application should be used as the initial query?, 2) which part(s) of the retrieved results should be used as the source of expansion and/or reduction?, 3) what is the best length of the expansion list in case of QE, or the best threshold for removing terms in case of QR?, 4) which pseudo-relevant results are really relevant and how many of them should be used?, and 5) what is the best relevance scoring model for the search task (e.g., BM25 [39], the vector space model with tf-idf weighting...etc).

Bouadjene et al. [27] provided a thorough evaluation on the CLEF-IP 2010/2011 collections to address some of the above challenges. The authors explored the scenario when only partial patent application is available for prior art search (e.g., *title*, *abstract*, *extended abstract*, or *description*). The authors tested different QE and QR general methods such as Rocchio [40] and a variant of the Maximal Marginal Relevance (MMR) [41]. They also tested patent-specific methods utilizing synonym sets [42], language models [29], and IPC-based lexicon [43]. After experimenting various sections as sources for the initial query terms as well as expansion/reduction sources, the results showed that, the *description* section among other sections is the best to use as the initial query in case of both QE and QR. QE was not beneficial for the long *description* queries as it already contains good coverage of relevant terms. However, QR on *description* queries was useful as it removed many of the

¹⁵This is a semi-supervised performance

¹⁶This is an Oracle performance

noisy terms. Generally, QR outperformed QE on *description* and *extended abstract* queries which indicates that, with long queries, QR is effective for better retrieval performance. The results also showed that generic QE methods such as Ricchio works generally better for QE than patent-specific QE methods. Finally, the results showed that BM25 scoring works better than the TF-IDF scoring on the long *description* queries for both QR and QE, while TF-IDF works better than BM25 on short and medium-length *title* or *abstract* queries. Through this comprehensive experimental study, the authors did not evaluate the impact of using multiple sections in combination as sources for QE or QR. More importantly, the study does not provide any insights about the respective values of number of expansion terms or term removal threshold and whether these values are somewhat deterministic or vary widely calling for interactive setting.

To address the problem of poor PRF results in PR compared to traditional IR, Bashir and Rauber [44] proposed a novel approach for PRF-based QE which builds a model that learns to identify better PRF results based on their similarity with the query patent over specific terms. These terms are learned by building a classification model that classifies whether a term would be useful for QE or not according to some proximity features between the original query terms and pseudo-relevant terms. The authors, through experiments on a subset of USPTO patents, showed the ability of this model to introduce more relevant QE terms and subsequently increasing the retrievability of individual patents. However, the authors did not evaluate this model on any of the available test collections. Moreover, extracting similarity features and computing similarities with PRF results during query execution is computationally expensive and time consuming.

Along the same efforts, Mahdabi and Crestani [28] proposed a framework for identifying effective PRF documents at runtime and then performing QE using terms from these relevant documents. The authors first proposed patent-specific features and then used them to build a regression model which calculates a relevancy score of each PRF document. Though results on the CLEF-IP 2010 collection were encouraging, several challenges still exist. For example, the computational complexity of calculating the regression model features at runtime. And PRF parameters tuning (e.g., number of PRF documents to use).

In [29], Ganguly et al. proposed a PRF approach which utilizes a language model for QR of long queries composed of full patent applications. The authors argued that, naive application of PRF to expand query terms could add noisy terms causing query-topic drift. Moreover, naive removal of terms that has unit term frequency in the query could cause removal of useful terms and thus hurt retrieval effectiveness. Instead, the authors proposed a PRF-based QR technique which generates language model similarity scores between query segments (sentences or n-grams) and top ranked results. Segments with top scores are kept and all others are removed. Results on the English subset of the CLEF-IP 2010 collection showed that the proposed approach outperforms the baselines. Parameter tuning is still the main downside of this technique. The performance of the proposed approach was unstable compared to the baselines with different parameter values. Specifically, the window size, the number of pseudo-relevant documents, and the fraction of terms to retain.

In [30], Golestan Far et al. provided a study on hybrid QRE which aims to automatically approximate the optimal \bar{Q} by careful selec-

tion/expansion of relevant query terms. To motivate the efficacy of QRE on retrieval performance, the authors first designed an experiment where relevance judgments of a query patent Q were assumed to be known in advance. After running Q , using PRF on top- K documents, only terms that are more frequent in retrieved relevant documents (those from relevance judgments) than irrelevant documents are kept and used as \bar{Q} . Then, querying using \bar{Q} achieved a better performance than state-of-the-art on the English subset of CLEF-IP 2010 collection. To approximate \bar{Q} automatically, the authors proposed four different methods hoping to identify relevant vs. irrelevant terms in Q by: 1) removing terms with high document frequency in the top-100 retrieved documents, 2) removing infrequent terms in Q , 3) using frequent terms in relevant documents assuming the top-5 retrieved documents are relevant, and 4) performing QR on Q using IPC definitions as stop-words. All of the four methods failed to perform better than the keyword-based baseline. More interestingly, the authors demonstrated that, baseline performance can be doubled if only one relevant document was manually provided by the user. This last observation motivates the need for interactive QRE as a simple and effective method for PR.

5.2.3 Semantic-based Methods

As we mentioned before, in PR queries can vary from few terms (e.g., survey memo) to thousands of terms (e.g., full patent application). Straightforward keyword-based PR proved to be ineffective simply because of the vocabulary mismatch between query terms and relevant patents content. Magdy et al. [19] showed that, in the CLEF-IP 2009 collection, 12% of the relevant documents have no common words with the search topics. This motivates the need for novel approaches to bridge this vocabulary mismatch gap. Several semantic-based methods have been proposed in attempt to match queries with relevant documents based on their meanings rather than relying on keyword matches only. Table 5 shows some semantic-based methods along with their performance results on benchmark datasets.

Dictionary-based: semantic-based methods perform QRE by expanding the query to include other terms that have similar meanings to the original query terms. The first category of these methods are the dictionary-based techniques which use either generic [42], technical [51], or patent-specific dictionaries [46, 47, 49, 50, 52, 53] for QRE. Generic dictionaries could be existing lexical databases such as WordNet [54], while patent-specific dictionaries are lexical databases generated from patent-related data such as examiner's query logs. In either case, similar or related terms to the original query terms are retrieved from such dictionaries and used for QE.

Magdy and Jones [42] explored the use of WordNet for QE in PR on the CLEF-IP 2010 collection. Overall, adding synonyms and hyponyms for nouns and verbs in the original query increased the MAP score slightly, while decreased the PRES score significantly. Moreover, query execution time was increased considerably. The authors considered this a "negative" result. As the use of WordNet was proven to be effective in other retrieval tasks [55, 56], more experiments are needed to affirm the authors' conclusion. For example, investigating the impact of using synonyms only or hyponyms only, and expanding terms belonging to specific sections or ambiguous terms only.

Recently, more research was focused on utilizing domain-specific and technical dictionaries rather than WordNet. Examiners' query logs have been an important resource for building such technical thesauri. Tannebaum et al. [45–50] introduced an analysis of the

Table 5: Semantic-based Methods

Method	Description	Dataset	MAP	PRES
Magdy & Jones [42]	<ul style="list-style-type: none"> • use Wordnet synonyms and hyponyms for QE • slow processing time • no improvement 	CLEF-IP 2010	<ul style="list-style-type: none"> • 0.136 • 0.140* 	<ul style="list-style-type: none"> • 0.484 • 0.486*
Tannebaum & Andreas [45–50]	<ul style="list-style-type: none"> • mine query logs for synonyms, co-occurring, and proximity terms • no improvement • use upon request 	CLEF-IP 2010	<ul style="list-style-type: none"> • 139 • 139* 	<ul style="list-style-type: none"> • 512 • 512*
Magdy & Jones [42]	<ul style="list-style-type: none"> • using synonyms learned from parallel translations (EN, GE, and FR) • improve MAP only • use upon request 	CLEF-IP 2010	<ul style="list-style-type: none"> • 144 • 140* 	<ul style="list-style-type: none"> • 485 • 486*

* indicates baseline performance

USPTO examiners’ search query logs. Their analysis, though on a subset of query logs, revealed interesting insights about patent examiners’ search behavior which could be very useful for designing effective PR systems. For example, the authors noted that about examiners’ behavior while searching for prior art: 1) the average query length is four terms, 2) search terms are mostly from the patent application under investigation, 3) expansion terms represent small percentage of query terms and mostly appear in the specific patent domain terminology, 4) the majority of query terms represents subject technical features that appears in the *claims* section, while very little percentage of them appears in the *description* section, 5) the majority of terms are nouns, followed by verbs, then adjectives, and 6) about half of the query operators used are "OR", followed by "AND", then proximity operators.

Tannebaum et al. built upon these insights and introduced methods to automatically identify synonyms/equivalents, co-occurring terms, and proximity relations for expanding query terms by mining examiners’ search logs. As we can notice, learning expansion terms from query logs might be misleading because not all query sessions succeed to identify prior art. Additionally, deeper analysis of the query logs considering other metadata such as relevant hits count might be useful in this regard. On the other hand, it would be more useful if we can model the features of these terms, for example, based on their location, frequency, part-of-speech...etc. From effectiveness perspective, evaluating the generated lexical knowledge on the CLEF-IP 2010 collection did not record significant improvement [49]. Therefore, the authors recommended using it in an interactive mode rather than automatic mode to semi-automate query generation.

Corpus-based: the second category of semantic-based QRE is the corpus-based methods. In these methods, textual corpora are analyzed to extract semantically related concepts to query terms which can be used for QE. Al-Shboul and Myaeng [57] proposed a Wikipedia-based QE method which works by first creating a summary of each Wikipedia article containing the main category, all titles under the main category, and other categories with in/out links to the main category. At query time, query terms and phrases are matched with page summaries, then, phrases from matching pages are scored and selected for QE under the assumption that they are semantically related. Experiments on the subset of USPTO patents in the NTCIR-6 collection showed increase in MAP over other QE techniques. However, the authors used IPC codes rather than citations as relevance judgments to topic queries which does not reflect

the typical search practices, where it is needed to retrieve related patent documents not related classification codes.

Another corpus-based method was proposed by Magdy and Jones [42], where synonym sets were automatically generated from the CLEF-IP patent corpus. The authors utilized parallel translations of patent sections in different languages to build a word-to-word translation model and infer synonymy relation when a word in one language is translated to multiple words in another language. These multiple words under some probabilistic threshold could be considered synonyms. Overall results using this method were better than PRF and Wordnet based QE, but worse than the keyword-based baseline in [36]. The authors also showed that, the performance of this method on some topics was better than the baseline which indicates its potential. The issue they raised is how to more effectively apply QE by selecting "good" terms [26], or predicting QE performance beforehand [58, 59]. Such challenges can also be alleviated semi-automatically by developing intelligent and usable interactive QE frameworks which engage users in such decision. Finally, Krestel and Smyth [60] applied topic modeling of search hits in order to better rank retrieved patents. The results on a small collection of the USPTO patents showed improved MAP.

5.2.4 Metadata-based Methods

Patents are not only textual documents, they contain lots of non-textual metadata and bibliographic information as well (e.g., citations, tables, formulas, drawings, classification...etc). Combining metadata analysis with text-based PR has shown improvements in performance in the literature [51, 61, 62, 64, 65]. Metadata features are also language independent making them advantageous when used for CLIR. Table 6 shows some metadata-based methods along with their performance results on benchmark datasets

Citation-based: The use of citation analysis for better PR is the most heavily reported technique of metadata-based methods. Naively incorporating citations from topic patent applications as prior art proved to be effective, eliminating the need for deeper citation analysis [66]. However, citation extraction from patent texts is challenging because there is no standard writing style for patent references. Lopez and Romary [64] developed a tool for citation mining which identifies, parses, normalizes, and consolidates patent citations. As citations might not be always available in all scenarios (e.g., related work search, technology survey...etc), more mature techniques are needed. Fujii [61] proposed using PageRank [67] and document popularity as an additional scoring to re-rank query

Table 6: Metadata-based Methods

Method	Description	Dataset	MAP	PRES
Fujii [61]	<ul style="list-style-type: none"> • use PageRank on patents citation graph • use patent popularity among top results with weighted voting 	NTCIR-6	<ul style="list-style-type: none"> • 0.075 • 0.081 • 0.071* 	-
Mahdabi & Crestani [62]	<ul style="list-style-type: none"> • build query specific citation graph from PRF results and their citations • weight nodes using PageRank • estimate query LM from the graph nodes considering their PageRank scores 	CLEF-IP 2011	<ul style="list-style-type: none"> • 0.105 • 0.099* 	<ul style="list-style-type: none"> • 0.481 • 0.450*
Mahdabi & Crestani [63]	<ul style="list-style-type: none"> • using time-aware random walk on weighted citation graph 	CLEF-IP 2011	<ul style="list-style-type: none"> • 0.125 • 0.058* 	<ul style="list-style-type: none"> • 0.536 • -

★ indicates baseline performance

top results returned using *claims*-based queries. The results of applying popularity scoring on the English subset of NTCIR-6 improved MAP and recall over the raw text-based scoring. Incorporating PageRank, though intuitive, poses many challenges especially because patent documents have references to non-patent literature which would produce incomplete citation graph. Mahdabi and Crestani [62] extended their query modeling technique in [20] by incorporating term distributions of the PRF results as well as their citations in calculating the query language model. The authors first construct a query-specific citation graph using PRF results and their citations and assign a score for each of them using PageRank. Then, a query model is estimated from term distributions of the documents in the citation graph constrained by their respective PageRank. Finally, QE is performed using the estimated query model. Experiments on the CLEF-IP 2011 collection showed improved recall performance with no change in precision, which indicates the usefulness of using cited documents vocabulary for QE. Best improvements were achieved using the top 30 PRF documents, 2-levels citation graph, and 100 expansion terms. However, we can notice two main computational challenges using this technique in real-time setting: 1) computing the PageRank of the 2-level citations graph, and 2) estimating the query model from top PRF documents as well as documents in the citation graph.

Classification-based: these methods utilize classification information of the topic patent and the retrieved documents to improve the performance of PR [68–72]. The naive use of IPC classification is to filter retrieved documents to keep only ones that share the same IPC classification code at some level (e.g., same subclass) with the topic patent [19, 73]. More sophisticated use of classification information was introduced by Verma and Varma in [74] who proposed a new representation of patent documents based on IPC classifications. The method utilizes IPC codes assigned to the corpus patents as well as codes of their citing documents to form an IPC class vector. First, the vector is initialized from patent’s IPC code, then codes of citing patents are propagated over multiple iterations. The most similar patents are retrieved using cosine similarity between IPC class vectors and re-ranked using text-based search utilizing the top 20 tf-idf topic patent terms. Experiments on the CLEF-IP 2011 collection showed improved recall but low MAP scores. The instability of the patent classification system poses a real challenge when it comes to incorporating classification metadata into PR systems. Overtime, new classes are added to the classification hierarchy and existing classes are expanded. In order to do reliable search based on classification codes, these changes must be accounted for periodically. Moreover, patents are assigned to multiple classifica-

tion codes, however, almost all previous research considered only the primary class but not secondary classifications which might, if utilized, improve the retrieval performance.

Hybrid: these methods utilize various sources of metadata to improve PR performance. In [63], Mahdabi and Crestani built upon previous work in [20, 62] and proposed a QE method that utilizes time-aware random walk on a weighted patent citations network. Citation weights are derived from various metadata (e.g., classification codes, inventors, assignee...etc). Citations with higher weights are considered more influential when performing QE. Experiments on the CLEF-IP 2010/2011 collections show improved recall and MAP. The authors in [53] proposed building a query-specific lexicon from IPC definition pages and using it for QE. Unfortunately, the lexicon would be helpful only if the query represents a complete patent document with IPC codes assigned to it which is not always the case especially at the early stages of the patent life-cycle.

5.2.5 Interactive Methods

Interactive PR is inevitable. As we can notice from the above review, effective fully automated retrieval of patent prior art is very challenging. Best methods perform around average in terms of PRES and much less in terms of MAP. Additionally, these methods require tuning a large number of parameters and thresholds whose optimal values differ according to the given query and the specific information need. For example, deciding which patent section to use?, which PRF results?, and which expansion terms and their respective weights?. The answers of these questions are not deterministic and probably require multiple interaction cycles with the user in order to satisfy his/her information need.

Current interactive methods in PR are more focused on better organization, integration, and utilization of structured and textual patent data than on better retrieval performance. In other words, PR is addressed as a professional search problem rather than prior art search problem. In [75], Fafalios and Tzitzikas presented a keyword-based interactive search framework to support patent search. The interaction elements are presented through post-analysis of search results in the form of facets based features like static metadata (e.g., IPC codes), textual clustering, named entity extraction, semantic enrichments, and others. The framework was applied on patent search [76] and evaluated using user study of twelve patent examiners [77]. Evaluation responses indicated overall acceptance of the framework in terms of usability, ease of use, efficiency, learnability. However, the authors did not report on the effectiveness or success of the system helping patent examiners to find prior art.




Figure 2: Concept graph of Cognitive Analytics. Light blue nodes are explicit concepts and red nodes are latent concepts.

Shalaby et al. [78] proposed a visual interactive semantic framework for patent analysis which features semantic-based QE of search queries using Mined Semantic Analysis (MSA) [79]. In a nutshell, MSA builds an association knowledge graph using rule mining of concept rich textual corpora (e.g., Wikipedia). After mining the "See Also" link graph of Wikipedia, MSA could represent a topic query as a bag-of-concepts (BOC) derived from the association knowledge graph. This BOC could then be used to expand the original query terms. Figure 2 shows an example of the query expansion map of *Cognitive Analytics*. Another example is presented in Figure 3 showing concept map of 10 patents of *Bank of America* using the *abstract* section. Users can interact with the concept map by removing nodes and updating the search results. Shalaby and Zadrozny [80] demonstrated the applicability of their framework to support tasks such as prior art search, competitive intelligence, technology landscape analysis and exploration. However, they did not provide a controlled study evaluating the performance of their method on benchmark collections.

Developing interactive methods for PR is also motivated by recent analysis which showed significant performance improvement if only one relevant document was manually provided by the user [30]. Performance gains using Technology Assisted Review (TAR) [81, 82] in domains like electronic discovery motivates investigating the applicability of machine learning TAR protocols in PR.

TAR, like PR, is a total-recall task where it is required to find all relevant documents to the search request with reasonable effort (time and cost). TAR is thus a human-in-the-loop process where a human expert manually annotates a subset of the documents as relevant or irrelevant. The underlying algorithm subsequently builds a ranking model by training on such annotations and uses this model to promote more relevant results and demote irrelevant ones as more documents are searched and annotated. This process stops when enough results are obtained. Typically, these algorithms utilize techniques such as continuous active learning combined with Boolean search in order to develop and adapt the ranking model [81].

Several questions still need to be addressed when it comes to investigating TAR protocols applicability to PR, as these protocols were only evaluated in ad hoc search scenarios. The complexity of patents terminology and availability of multiple sources of meta-data would, likely, demonstrate many opportunities for adaptation and modifications to the current TAR protocols.

6. RELATED TOPICS

Despite the intense interest within the research community in PR, the patent industry has many other challenges and open problems which are of high interest and value to various stakeholders such as

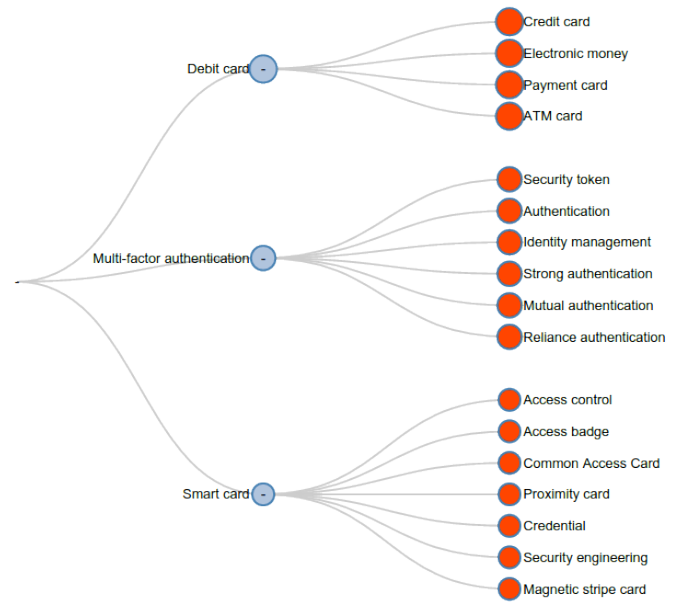


Figure 3: Concept graph using Bank of America's 100 patent titles. Light blue nodes are explicit concepts and red nodes are latent concepts.

economists, R&D managers, and legal professionals to name a few. In this section, we try to lightly touch on these tasks and highlight some challenges and possible future directions.

6.1 Patent Quality Assessment

assessing the technical quality and importance of inventions is very important to patent owners because it allows them to:

- better utilize their IP management costs by automated recommendation of patent maintenance decisions.
- better determine the novelty and originality of their patents.
- maximize licensing revenues by automatic estimation of the patent value.

Because there is no ground truth for quality measurements, performance evaluation of quality assessment techniques is usually based on indicators such as correlation with patent forward citations, maintenance status history, court rulings (if any), and/or patent reexamination history (if any). Some early work scored patents using their metadata such as citations count, maintenance history, global prosecution efforts [83], and even manually by patent attorneys. Automated patent quality assessment has gained more traction in recent years though.

Citation analysis has been and still a main technique for patent valuation [84–88]. Wang et al. [88] proposed a probabilistic mixture approach to predict whether a topic patent will be renewed at different renewal periods. The method first divides the citations into two groups; technological and legal. From each group, different features reflecting the technological richness, technological influence, legal patent scope, and legal blocking power of each patent are combined. The authors subsequently build a binary classifier using

these probabilistic features. Evaluation is performed by comparing the model's predictions against the renewal decisions of a collection of patents. While proved effectiveness, estimating patent value as a binary outcome might not be practical especially if a patent owner needs to prioritize his maintenance decisions of multiple patents.

Quality assessment based on the lexical features of the patent text was also explored in the literature [89–91]. In [90], Liu et al. proposed a graphical model to estimate patent quality as a latent variable. The model utilized lexical features extracted from the patent text such as *claims* n-grams age and popularity, lexical alignment between the *claims* and the *description*, number of dependent and independent claims, number of reported classes when filing the patent, and other features. The authors also incorporated measurements such as forward citations count, court decisions, and reexamination records. It is clear that court decisions are only available for small number of patents which might not allow building a robust model.

Jin et al. [89] modeled the patent maintenance decision as recommendation problem where patents were represented as multi-modal heterogeneous information network. The model utilized several metadata features, lexical features such as unique words and lengths of different sections, as well as inventor and assignee profile features. Experimental results showed high prediction accuracy on a large number of USPTO patents.

Hu et al. [92] proposed a time-based topic model which ranks patents novelty and influence based on whether the dominant topics in patent's prior art (for novelty) or forward art (for influence) are still active topics. The authors also proposed using time decay function to address the problem of old patents having less prior art and more forward art than newer patents and vice versa. Results showed high correlation between assigned ranks and forward citations count.

Hido et al. [91] proposed a scoring model which assigned a patentability score to each patent and thus can be utilized to determine whether it will be granted. First, the authors extracted textual features such as word frequency, word age, and syntactic complexity (e.g., number of sentences). Then, they trained a classifier using previous patent office decisions as ground truth. Though results showed the model effectiveness, the utilized syntactic complexity features are all extracted from the topic patent and thus could be good predictors for the writing quality not patentability potential.

The correlation between patent *claims* novelty and patent value has been analyzed in [93, 94] using lexical analysis of patent text. In [94], Hasan et al. proposed an IR-based ranking tool which analyses patent *claims* for originality. The technique first extracts key terms and phrases from the *claims* text using syntactic patterns and then looks for usage patterns backward to determine their novelty, and forward to determine their influence. The method considers usage patterns only through user defined time window. It is also keyword-based and hence will fail to capture key phrases that are semantically similar and subsequently might give inaccurate scores.

Along the efforts of using patent legal data for quality assessment Mann and Underweiser [95] utilized prosecution histories, court decisions, and patent textual features to analyze patent quality. The analysis suggested that patent examination records would be very helpful in better discriminating high from low quality patents and possibly improve the examination process as a whole.

6.2 Patent Litigation

Litigation in general, and patent litigation specifically have been and still a topic of interest to legal professionals. With the increased amounts of digitized data available and the need for technology support in analyzing and mining these huge datasets, litigation became of more interest to computational science researchers. Patent litigation can take many forms, the most common is patent infringement litigation where a patent owner (plaintiff) accuses another party (defendant) of using his/her invention without license or permission. Because litigation is very expensive, the most common defensive action for the defendant is to establish invalidity of the plaintiff invention by issuing a post grant proceeding such as PGR or IPR. Now the problem becomes a PR task, i.e. invalidity search, where one of the aforementioned methods can be utilized with wider scope to cover not only patent literature but also other published material.

The task of automatically establishing patent infringement is not addressed in literature. Such task requires extensive human expertise and reasoning to build correspondences between product features and patent claims. On the other hand, statistical and visual analytics of previous court decisions have shown some degree of success in helping lawyers to better understand possible outcomes and better plan on defense strategies [96–98].

For example, Allison et al. [99] provided a statistical study on patent cases filed from 2008 to 2009 and decisions made between (2009-2013). The study showed that, there is a strong correlation between court decision, and patent-specific, litigation-specific, and industry-specific variables such as industry and technology type, inventors foreign status, number of claims, number of forward and backward citations, and number of defendants sued.

Rajshekhar et al. [100] studied the potential of concept-based semantic search in patent litigation. The authors designed an experiment in order to retrieve invalidating patents to a given litigated patent using a subset of PTAB's final decisions as ground truth and a search corpus of ~7m USPTO patents. The authors, based on the experimental results and through interviews with patent practitioners, concluded that, a one-size-fits-all semantic search approach is incapable of capturing the highly nuanced relevance judgments made in the domain of patent litigation. Rather, the search workflow should be modeled as a multistage information seeking process, where users are presented with interactive elements to control the search space, and their feedback is incorporated iteratively in the relevance ranking of retrieved results for enhanced performance.

Finally, There is much to be done in building predictive models for patent litigation given the availability of prior case datasets that were not available few years ago (e.g., prosecution histories, court decisions, and PTAB decisions).

6.3 Technology Licensing

Patents represent one of the most valuable assets in today's enterprises which, if leveraged effectively, guarantee not only competitive superiority, but also huge licensing revenues [93]. The technology licensing task is three sided. First, patent owners would be interested in finding potential licensees with reasonable effort. Second, licensees would like to find relevant inventions to their businesses. Third, owners and businesses would be interested in gauging the strategic and protection values of a patent in order to support their pricing and offering decisions.

While there is no much research focusing on the task of automatically recommending potential licensees. The task of recommending patents to be licensed was relatively more considered. Chen et al. [93] proposed a platform called SIMPLE which is used at IBM to identify target patents for licensing. Given a set of topic patents, SIMPLE uses nearest neighbor similarity to find other patents that are most similar to the given topic set. Then, all the patents are grouped and proposed as one licensing package to interested party. The platform was extended in [101] to allow retrieving target patents using free text search. We can notice that current trends for identifying potential patents for licensing model the problem as a PR task. More elaboration on the SIMPLE platform was introduced by Spangler et al. [102] using interactive visualization. First, portfolios of two companies are contrasted to find content overlap between both of them using proximal search. Then, the closest patents to the overlap area are recommended as candidates for licensing.

7. CONCLUDING REMARKS

In this paper we presented a comprehensive review of PR methods and approaches. It is clear that, the well-performing IR techniques in areas like Web search can not be utilized directly in PR without deliberate domain adaptation and customization. Furthermore, state-of-the-art performance in automatic PR is still low (<0.2 MAP). Several proposed techniques for query expansion, query reduction and pseudo relevance feedback require tuning of various parameters. Search professional practices suggest that effective prior art search requires multiple iterations of searching, reviewing, and refining. On the other hand, examiners' query formulation practices (few keywords and Boolean search) are different from those of automatic methods (many keywords and free-text search). These observations motivate the need for interactive search tools which provide cognitive assistance to search professionals with minimal effort. These tools must also be developed in hand with patent professionals considering their practices and expectations.

Unexplored patent-related data sources might be an opportunity for breakthrough improvements over the current modest state-of-the-art in PR. For example, utilizing reexamination records, PTAB decisions, differences between the patent application and the granted version, examiner/applicant correspondences, and prosecution histories. All these resources are not yet fully explored in the literature of PR.

Related tasks such as patent quality assessment, litigation, and licensing are of less focus among computational scientists. However, they provide wide opportunities for future exploration from computational and modeling perspectives. These tasks require interdisciplinary and cooperative efforts from both legal professionals and the computer science research community.

8. REFERENCES

- [1] F. Piroi, M. Lupu, A. Hanbury, and V. Zenz, "Clef-ip 2011: Retrieval in the intellectual property domain.," in CLEF (notebook papers/labs/workshop), Citeseer, 2011.
- [2] M. Lupu, K. Mayer, J. Tait, and A. J. Trippe, Current challenges in patent information retrieval, vol. 29. Springer Science & Business Media, 2011.
- [3] F. Piroi, M. Lupu, and A. Hanbury, Information Access Evaluation. Multilinguality, Multimodality, and Visualization: 4th International Conference of the CLEF Initiative, CLEF 2013, Valencia, Spain, September 23-26, 2013. Proceedings, ch. Overview of CLEF-IP 2013 Lab, pp. 232–249. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013.
- [4] G. Roda, J. Tait, F. Piroi, and V. Zenz, Multilingual Information Access Evaluation I. Text Retrieval Experiments: 10th Workshop of the Cross-Language Evaluation Forum, CLEF 2009, Corfu, Greece, September 30 - October 2, 2009, Revised Selected Papers, ch. CLEF-IP 2009: Retrieval Experiments in the Intellectual Property Domain, pp. 385–409. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010.
- [5] F. Piroi, "Clef-ip 2010: Retrieval experiments in the intellectual property domain,"
- [6] F. Piroi, M. Lupu, A. Hanbury, W. Magdy, A. Sexton, and I. Filippov, CLEF-IP 2012: Retrieval experiments in the intellectual property domain, vol. 1178. CEUR-WS, 2012.
- [7] E. Graf and L. Azzopardi, "A methodology for building a patent test collection for prior art search," in In Proceedings of the 2nd International Workshop on Evaluating Information Access, EVIA, 2008.
- [8] NTCIR, "NTCIR Test Collections," 2015. [<http://research.nii.ac.jp/ntcir/permission/data-en.htm>; accessed 30-April-2016].
- [9] M. Iwayama, A. Fujii, N. Kando, and A. Takano, "Overview of patent retrieval task at ntcir-3," in Proceedings of the ACL-2003 workshop on Patent corpus processing-Volume 20, pp. 24–32, Association for Computational Linguistics, 2003.
- [10] A. Fujii, M. Iwayama, and N. Kando, "Overview of patent retrieval task at ntcir-4.," in NTCIR, 2004.
- [11] A. Fujii, M. Iwayama, and N. K., "Overview of patent retrieval task at ntcir-5," in In Proceedings of the Fifth NTCIR Workshop Meeting on Evaluation of Information Access Technologies: Information Retrieval, Question Answering and Cross-lingual Information Access, pp. 269–277, 2005.
- [12] A. Fujii, M. Iwayama, and N. Kando, "Overview of the patent retrieval task at the ntcir-6 workshop.," in NTCIR, 2007.
- [13] M. Lupu, J. Huang, J. Zhu, and J. Tait, "Trec-chem: large scale chemical information retrieval evaluation at trec.," in ACM SIGIR Forum, vol. 43, pp. 63–70, ACM, 2009.
- [14] M. Lupu, J. Tait, J. Huang, and J. Zhu, "Trec-chem 2010: Notebook report," Proceedings of TREC 2010, vol. 2, 2010.
- [15] M. Lupu, J. Zhao, J. Huang, H. Gurulingappa, J. Fluck, M. Zimmermann, I. V. Filippov, and J. Tait, "Overview of the trec 2011 chemical ir track.," in TREC, 2011.
- [16] D. L. Schwartz and T. M. Sichelman, "Data sources on patents, copyrights, trademarks, and other intellectual property," Copyrights, Trademarks, and Other Intellectual Property (August 17, 2015), vol. 2, 2015.
- [17] J. J. Jürgens, P. Hansen, and C. Womser-Hacker, "Going beyond clef-ip: The reality for patent searchers?," in Information Access Evaluation. Multilinguality, Multimodality, and Visual Analytics, pp. 30–35, Springer, 2012.
- [18] S. Verberne and E. D'hondt, "Prior art retrieval using the claims section as a bag of words," in Multilingual Information Access Evaluation I. Text Retrieval Experiments, pp. 497–501, Springer, 2009.
- [19] W. Magdy, J. Leveling, and G. J. Jones, "Exploring structured documents and query formulation techniques for

- patent retrieval,” in Multilingual Information Access Evaluation I. Text Retrieval Experiments, pp. 410–417, Springer, 2009.
- [20] P. Mahdabi, M. Keikha, S. Gerani, M. Landoni, and F. Crestani, Building queries for prior-art search. Springer, 2011.
- [21] F. Wang, L. Lin, S. Yang, and X. Zhu, “A semantic query expansion-based patent retrieval approach,” in Fuzzy Systems and Knowledge Discovery (FSKD), 2013 10th International Conference on, pp. 572–577, IEEE, 2013.
- [22] K. Konishi, “Query terms extraction from patent document for invalidity search,” in NTCIR, 2005.
- [23] R. Baeza-Yates, B. Ribeiro-Neto, et al., Modern information retrieval, vol. 463. ACM press New York, 1999.
- [24] K. Järvelin and J. Kekäläinen, “Cumulated gain-based evaluation of ir techniques,” ACM Transactions on Information Systems (TOIS), vol. 20, no. 4, pp. 422–446, 2002.
- [25] W. Magdy and G. J. Jones, “Pres: a score metric for evaluating recall-oriented information retrieval applications,” in Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval, pp. 611–618, ACM, 2010.
- [26] G. Cao, J.-Y. Nie, J. Gao, and S. Robertson, “Selecting good expansion terms for pseudo-relevance feedback,” in Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval, pp. 243–250, ACM, 2008.
- [27] M. R. Bouadjeneq, S. Sanner, and G. Ferraro, “A study of query reformulation for patent prior art search with partial patent applications,” in Proceedings of the 15th International Conference on Artificial Intelligence and Law, pp. 23–32, ACM, 2015.
- [28] P. Mahdabi and F. Crestani, “Learning-based pseudo-relevance feedback for patent retrieval,” in Multidisciplinary Information Retrieval, pp. 1–11, Springer, 2012.
- [29] D. Ganguly, J. Leveling, W. Magdy, and G. J. Jones, “Patent query reduction using pseudo relevance feedback,” in Proceedings of the 20th ACM international conference on Information and knowledge management, pp. 1953–1956, ACM, 2011.
- [30] M. Golestan Far, S. Sanne, M. R. Bouadjeneq, G. Ferraro, and D. Hawking, “On term selection techniques for patent prior art search,” in Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 803–806, ACM, 2015.
- [31] S. Verberne, “Clef-ip 2010: Prior art retrieval using the different sections in patent documents,”
- [32] M. Z. Wanagiri and M. Adriani, “Prior art retrieval using various patent document fields contents,”
- [33] E. Meij, W. Weerkamp, and M. de Rijke, “A query model based on normalized log-likelihood,” in Proceedings of the 18th ACM conference on Information and knowledge management, pp. 1903–1906, ACM, 2009.
- [34] D. Hiemstra, S. Robertson, and H. Zaragoza, “Parsimonious language models for information retrieval,” in Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval, pp. 178–185, ACM, 2004.
- [35] X. Xue and W. B. Croft, “Transforming patents into prior-art queries,” in Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval, pp. 808–809, ACM, 2009.
- [36] W. Magdy and G. J. Jones, “Applying the kiss principle for the clef-ip 2010 prior art candidate patent search task,” 2010.
- [37] M. Osborn, T. Strzalkowski, and M. Marinescu, “Evaluating document retrieval in patent database: a preliminary report,” in Proceedings of the sixth international conference on Information and knowledge management, pp. 216–221, ACM, 1997.
- [38] F. Wang and L. Lin, “Query construction based on concept importance for effective patent retrieval,” in Fuzzy Systems and Knowledge Discovery (FSKD), 2015 12th International Conference on, pp. 1455–1459, IEEE, 2015.
- [39] S. E. Robertson, S. Walker, S. Jones, M. M. Hancock-Beaulieu, M. Gatford, et al., “Okapi at trec-3,” NIST SPECIAL PUBLICATION SP, vol. 109, p. 109, 1995.
- [40] G. Salton, The SMART Retrieval System—Experiments in Automatic Document Processing. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1971.
- [41] J. Carbonell and J. Goldstein, “The use of mmr, diversity-based reranking for reordering documents and producing summaries,” in Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval, pp. 335–336, ACM, 1998.
- [42] W. Magdy and G. J. Jones, “A study on query expansion methods for patent retrieval,” in Proceedings of the 4th workshop on Patent information retrieval, pp. 19–24, ACM, 2011.
- [43] P. Mahdabi, S. Gerani, J. X. Huang, and F. Crestani, “Leveraging conceptual lexicon: query disambiguation using proximity information for patent retrieval,” in Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval, pp. 113–122, ACM, 2013.
- [44] S. Bashir and A. Rauber, “Improving retrievability of patents in prior-art search,” in Advances in Information Retrieval, pp. 457–470, Springer, 2010.
- [45] W. Tannebaum and A. Rauber, “Acquiring lexical knowledge from query logs for query expansion in patent searching,” in Semantic Computing (ICSC), 2012 IEEE Sixth International Conference on, pp. 336–338, IEEE, 2012.
- [46] W. Tannebaum and A. Rauber, “Analyzing query logs of uspto examiners to identify useful query terms in patent documents for query expansion in patent searching: a preliminary study,” in Multidisciplinary Information Retrieval, pp. 127–136, Springer, 2012.
- [47] W. Tannebaum and A. Rauber, “Mining query logs of uspto patent examiners,” in Information Access Evaluation. Multilinguality, Multimodality, and Visualization, pp. 136–142, Springer, 2013.
- [48] W. Tannebaum and A. Rauber, “Using query logs of uspto patent examiners for automatic query expansion in patent searching,” Information Retrieval, vol. 17, no. 5-6, pp. 452–470, 2014.
- [49] W. Tannebaum, P. Mahdabi, and A. Rauber, “Effect of log-based query term expansion on retrieval effectiveness in patent searching,” in Experimental IR Meets

Multilinguality, Multimodality, and Interaction, pp. 300–305, Springer, 2015.

- [50] W. Tannebaum and A. Rauber, “Patnet: a lexical database for the patent domain,” in Advances in Information Retrieval, pp. 550–555, Springer, 2015.
- [51] P. Lopez and L. Romary, “Multiple retrieval models and regression models for prior art search,” in CLEF 2009 Workshop, p. 18p, 2009.
- [52] J. Wajda and W. Zadrozny, Challenging Problems and Solutions in Intelligent Systems, ch. Prior-Art Relevance Ranking Based on the Examiner’s Query Log Content, pp. 323–333. Cham: Springer International Publishing, 2016.
- [53] P. Mahdabi and F. Crestani, “Patent query formulation by synthesizing multiple sources of relevance evidence,” ACM Transactions on Information Systems (TOIS), vol. 32, no. 4, p. 16, 2014.
- [54] C. Fellbaum, WordNet: An Electronic Lexical Database. Bradford Books, 1998.
- [55] E. M. Voorhees, “Using wordnet for text retrieval,” Fellbaum (Fellbaum, 1998), pp. 285–303, 1998.
- [56] S. Liu, F. Liu, C. Yu, and W. Meng, “An effective approach to document retrieval via utilizing wordnet and recognizing phrases,” in Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval, pp. 266–272, ACM, 2004.
- [57] B. Al-Shboul and S.-H. Myaeng, “Wikipedia-based query phrase expansion in patent class search,” Information Retrieval, vol. 17, no. 5-6, pp. 430–451, 2014.
- [58] S. Cronen-Townsend, Y. Zhou, and W. B. Croft, “Predicting query performance,” in Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval, pp. 299–306, ACM, 2002.
- [59] Y. Lv and C. Zhai, “Adaptive relevance feedback in information retrieval,” in Proceedings of the 18th ACM conference on Information and knowledge management, pp. 255–264, ACM, 2009.
- [60] R. Krestel and P. Smyth, “Recommending patents based on latent topics,” in Proceedings of the 7th ACM conference on Recommender systems, pp. 395–398, ACM, 2013.
- [61] A. Fujii, “Enhancing patent retrieval by citation analysis,” in Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval, pp. 793–794, ACM, 2007.
- [62] P. Mahdabi and F. Crestani, “The effect of citation analysis on query expansion for patent retrieval,” Information Retrieval, vol. 17, no. 5-6, pp. 412–429, 2014.
- [63] P. Mahdabi and F. Crestani, “Query-driven mining of citation networks for patent citation retrieval and recommendation,” in Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, pp. 1659–1668, ACM, 2014.
- [64] P. Lopez and L. Romary, “Experiments with citation mining and key-term extraction for prior art search,” in CLEF 2010-Conference on Multilingual and Multimodal Information Access Evaluation, 2010.
- [65] D. Eisinger, G. Tsatsaronis, M. Bundschuh, U. Wieneke, and M. Schroeder, “Automated patent categorization and guided patent search using ipc as inspired by mesh and pubmed,” Journal of biomedical semantics, vol. 4, no. 1, p. 1, 2013.
- [66] W. Magdy, P. Lopez, and G. J. Jones, “Simple vs. sophisticated approaches for patent prior-art search,” in Advances in Information Retrieval, pp. 725–728, Springer, 2011.
- [67] S. Brin and L. Page, “Reprint of: The anatomy of a large-scale hypertextual web search engine,” Computer networks, vol. 56, no. 18, pp. 3825–3833, 2012.
- [68] J. Kim, I.-S. Kang, and J.-H. Lee, “Cluster-based patent retrieval using international patent classification system,” in Computer Processing of Oriental Languages. Beyond the Orient: The Research Challenges Ahead, pp. 205–212, Springer, 2006.
- [69] C. G. Harris, S. Foster, R. Arens, and P. Srinivasan, “On the role of classification in patent invalidity searches,” in Proceedings of the 2nd international workshop on Patent information retrieval, pp. 29–32, ACM, 2009.
- [70] C. G. Harris, R. Arens, and P. Srinivasan, “Comparison of ipc and uspc classification systems in patent prior art searches,” in Proceedings of the 3rd international workshop on Patent Information Retrieval, pp. 27–32, ACM, 2010.
- [71] Y.-L. Chen and Y.-T. Chiu, “An ipc-based vector space model for patent retrieval,” Information Processing & Management, vol. 47, no. 3, pp. 309–322, 2011.
- [72] A. Giachanou, M. Salampasis, and G. Paltoglou, “Multilayer source selection as a tool for supporting patent search and classification,” Information Retrieval Journal, vol. 18, no. 6, pp. 559–585, 2015.
- [73] J. Gobeill, E. Pasche, D. Teodoro, and P. Ruch, “Simple pre and post processing strategies for patent searching in clef intellectual property track 2009,” in Multilingual Information Access Evaluation I. Text Retrieval Experiments, pp. 444–451, Springer, 2009.
- [74] M. Verma and V. Varma, “Patent search using ipc classification vectors,” in Proceedings of the 4th workshop on Patent information retrieval, pp. 9–12, ACM, 2011.
- [75] P. Fafalios and Y. Tzitzikas, “Exploratory professional search through semantic post-analysis of search results,” in Professional Search in the Modern World, pp. 166–192, Springer, 2014.
- [76] M. Salampasis and A. Hanbury, “Perfedpat: An integrated federated system for patent search,” World Patent Information, vol. 38, pp. 4–11, 2014.
- [77] M. Salampasis, A. Giachanou, and A. Hanbury, “An evaluation of an interactive federated patent search system,” in Multidisciplinary Information Retrieval, pp. 120–131, Springer, 2014.
- [78] W. Shalaby and W. Zadrozny, “Innovation analytics using mined semantic analysis,” in Proceedings of the 29th International FLAIRS Conference, 2016.
- [79] W. Shalaby and W. Zadrozny, “Measuring semantic relatedness using mined semantic analysis,” arXiv preprint arXiv:1512.03465, 2015.
- [80] W. Shalaby, K. Rajshekhar, and W. Zadrozny, “A visual semantic framework for innovation analytics,” in Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16).
- [81] M. R. Grossman and G. V. Cormack, “Technology-assisted review in e-discovery can be more effective and more efficient than exhaustive manual review,” Rich. JL & Tech., vol. 17, pp. 11–16, 2011.

- [82] G. V. Cormack and M. R. Grossman, "Evaluation of machine-learning protocols for technology-assisted review in electronic discovery," in Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval, pp. 153–162, ACM, 2014.
- [83] J. O. Lanjouw, A. Pakes, and J. Putnam, "How to count patents and value intellectual property: The uses of patent renewal and application data," The Journal of Industrial Economics, vol. 46, no. 4, pp. 405–432, 1998.
- [84] M. Trajtenberg, "A penny for your quotes: patent citations and the value of innovations," The Rand Journal of Economics, pp. 172–187, 1990.
- [85] D. Harhoff, F. Narin, F. M. Scherer, and K. Vopel, "Citation frequency and the value of patented inventions," Review of Economics and statistics, vol. 81, no. 3, pp. 511–515, 1999.
- [86] B. H. Hall, A. Jaffe, and M. Trajtenberg, "Market value and patent citations," RAND Journal of economics, pp. 16–38, 2005.
- [87] D. Czarnitzki, K. Hussinger, and B. Leten, "The market value of blocking patents," 2010.
- [88] S. Wang, Z. Lei, and W.-C. Lee, "Exploring legal patent citations for patent valuation," in Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, pp. 1379–1388, ACM, 2014.
- [89] X. Jin, S. Spangler, Y. Chen, K. Cai, R. Ma, L. Zhang, X. Wu, and J. Han, "Patent maintenance recommendation with patent information network model," in Data Mining (ICDM), 2011 IEEE 11th International Conference on, pp. 280–289, IEEE, 2011.
- [90] Y. Liu, P.-y. Hseuh, R. Lawrence, S. Meliksetian, C. Perlich, and A. Veen, "Latent graphical models for quantifying and predicting patent quality," in Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 1145–1153, ACM, 2011.
- [91] S. Hido, S. Suzuki, R. Nishiyama, T. Imamichi, R. Takahashi, T. Nasukawa, T. Idé, Y. Kanehira, R. Yohda, T. Ueno, et al., "Modeling patent quality: A system for large-scale patentability analysis using text mining," Information and Media Technologies, vol. 7, no. 3, pp. 1180–1191, 2012.
- [92] P. Hu, M. Huang, P. Xu, W. Li, A. K. Usadi, and X. Zhu, "Finding nuggets in ip portfolios: core patent mining through textual temporal analysis," in Proceedings of the 21st ACM international conference on Information and knowledge management, pp. 1819–1823, ACM, 2012.
- [93] Y. Chen, S. Spangler, J. Kreulen, S. Boyer, T. D. Griffin, A. Alba, A. Behal, B. He, L. Kato, A. Lelescu, et al., "Simple: a strategic information mining platform for licensing and execution," in Data Mining Workshops, 2009. ICDMW'09. IEEE International Conference on, pp. 270–275, IEEE, 2009.
- [94] M. A. Hasan, W. S. Spangler, T. Griffin, and A. Alba, "Coa: Finding novel patents through text analysis," in Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 1175–1184, ACM, 2009.
- [95] R. J. Mann and M. Underweiser, "A new look at patent quality: Relating patent prosecution to validity," Journal of Empirical Legal Studies, vol. 9, no. 1, pp. 1–32, 2012.
- [96] T. Harbert, "The law machine," Spectrum, IEEE, vol. 50, no. 11, pp. 31–54, 2013.
- [97] J. R. Allison, M. A. Lemley, and D. L. Schwartz, "Understanding the realities of modern patent litigation," 2014.
- [98] M. K. Osbeck, "Using data analytics tools to supplement traditional research and analysis in forecasting case outcomes," U of Michigan Public Law Research Paper Series, no. 446, 2015.
- [99] J. R. Allison, M. A. Lemley, and D. L. Schwartz, "Our divided patent system," The University of Chicago Law Review, pp. 1073–1154, 2015.
- [100] K. Rajshekhar, W. Shalaby, and W. Zadrozny, "Analytics in post-grant patent review: Possibilities and challenges (preliminary report)," in Proceedings of the American Society for Engineering Management 2016 International Annual Conference, 2016.
- [101] S. Spangler, Y. Chen, J. Kreulen, S. Boyer, T. Griffin, A. Alba, L. Kato, A. Lelescu, and S. Yan, "Simple: Interactive analytics on patent data," in Data Mining Workshops (ICDMW), 2010 IEEE International Conference on, pp. 426–433, IEEE, 2010.
- [102] S. Spangler, C. Ying, J. Kreulen, S. Boyer, T. Griffin, A. Alba, L. Kato, A. Lelescu, and S. Yan, "Exploratory analytics on patent data sets using the simple platform," World Patent Information, vol. 33, no. 4, pp. 328–339, 2011.